

Round Up The Usual Suspects: Knowledge-Based Metaphor Generation

Tony Veale

School of Computer Science
University College Dublin
Belfield, Dublin D4, Ireland.

Tony.Veale@UCD.ie

Abstract

The elasticity of metaphor as a communication strategy has spurred philosophers to question its ability to mean anything at all. If a metaphor can elicit different responses from different people in varying contexts, how can one say it has a *single* meaning? Davidson has argued that metaphors have no *special* or *secondary* meaning, and must thus mean exactly what they seem to mean on the surface. It is this literally anomalous meaning that directs us to the pragmatic inferences that a speaker actually wishes us to explore. Conveniently, this *laissez faire* strategy assumes that metaphors are crafted from apt knowledge by speakers with real communicative intent, allowing useful inference to be extracted from their words. But this assumption is not valid in the case of many machine-generated metaphors that merely echo the linguistic form – but lack the actual substance – of real metaphors. We present here an open public resource with which a metaphor-generation system can give its figurative efforts real meaning.

1 The Dreamwork of Language

Metaphor is the rubber cement of language. Not only does it help us to the plug holes on our lexica, we also use it to fill the gaps in our understanding and to hide the cracks in our arguments. For unlike the brittle plaster of literal language, metaphors are elastic and can readily expand to fit our meanings in a shifting conversational context. This elasticity comes at a price, though one which a master orator is happy to pay: our metaphors are elastic because they are indeterminate, underspecified and vague.

Like dreams, our metaphors paint vivid pictures with words, albeit with fuzzy and ill-defined edges. Like dreams, metaphors can be highly suggestive, yet leave us feeling confused and uncertain.

If metaphorical images are crisp at their focal points but hazy and dreamlike at their edges, just what is the meaning of any metaphor? The philosopher Donald Davidson (1978) has controversially argued that, like our dreams, our metaphors do not have well-defined meanings, at least not of a kind that an AI researcher, semanticist or computational linguist could squeeze into a symbolic structure. Rather, metaphors can move us to think and feel in certain ways, and perhaps act in certain ways, but like dreams, two analysts (a Jungian and a Freudian, say) can hold conflicting views as to how they should be interpreted and as to what they actually “mean”, if anything. So, for Davidson, a metaphor means just what it purports to mean on the surface, that is, what the literal or dictionary senses of its words suggest that it means. This *meaning* is to be distinguished from the panoply of inferences and insights that might later emerge from a metaphor, for regardless of how salient these may seem, they *cannot* be considered its definitive meaning. Freud once joked that when it comes to dreams, a cigar is often just a cigar. For Davidson, a figurative cigar is *always* a cigar, even if the metaphor spurs us to further inference far beyond the realm of tobacco.

If all metaphors mean simply what they seem to mean on the surface, and most – from the very best to the truly awful – are superficially anomalous, how can we tell the good from the bad by simply looking? Indeed, how can we tell *real* metaphors from *fake* metaphors based only on the words they use and their senses in the dictionary? Empirical results seem to bear out Davidson’s intuitions regarding our folk grasp of metaphors. Veale (2015)

reports on experiments in which the outputs of two automated generators of metaphors are compared and contrasted. The first, *@metaphorminute* (from noted Twitterbot builder Darius Kazemi) simply fills the copula template “*A is a B: X and Y*” with more-or-less random word choices for the tenor A, the vehicle B and putative grounds X and Y. The second, *@MetaphorMagnet* (Veale *et al.*, 2015) uses a knowledge-base of stereotypical properties and norms to generate comparisons grounded in a high-degree of similarity *and* semantic tension. In a crowd-sourced evaluation of both, test subjects were asked to rate the perceived comprehensibility of metaphors sampled from each generator. While 75% of *@MetaphorMagnet*’s outputs were said to have medium-high to very-high understandability, over half of *@metaphorminute*’s random offerings were rated just as comprehensible. In fact, because the latter were unconstrained by any representation of knowledge or meaning, subjects rated them as *more* novel than the knowledge-driven metaphors of *@MetaphorMagnet*. Only when subjects were asked to complete partial metaphors in a *cloze* test, to show a real understanding of the author’s intent, was the impact of knowledge, meaning and intent made clear: test subjects were able to reconstitute *@MetaphorMagnet*’s outputs with medium-high to very-high confidence in 78% of cases, but were unable to complete the original *@metaphorminute* metaphors with better than random performance.

Davidson’s *laissez-faire* take on meaning in metaphor can only take machines so far, and if we want a generator to spark the expected inferences in readers, we must assume a computational notion of meaning that goes deeper than the literal surface. For if we want our automatic generators to achieve more than the mere appearance of comprehensibility and forcefully make their points via metaphor, we must first give them knowledge of the world in which to anchor the insights conveyed by their metaphors. While this is a *big ask* that cannot be realized with any single knowledge-base, we present here a solid open-source foundation for the development of a figurative knowledge system. This knowledge-base, named *the NOC List (Non-Official Characterization list)* is a source of stereotypical knowledge regarding popular culture, famous people (real and fictional) and their trademark qualities, behaviours and settings. We show here how this modular resource, designed to grow and evolve, can be used for metaphor generation.

2 The Representational Core of Metaphor

Despite Davidson’s withering views on the *hidden* meaning of metaphor, and also, in part, *because* of his radical skepticism, it seems necessary for an AI system to find common ground between the literal or surface meaning of a metaphor and its eventual interpretation in the mind of a listener, since an understanding of one is needed for an appreciation of the other. As such, most symbolic approaches to metaphor interpretation assume that these surface and deep meanings *must* overlap in some representational form, for it is this overlap that allows an NLP system to find its way from one to the other.

Many systems assume that this overlap is *taxonomic* in nature: simply, that the concepts used in the surface meaning share a general categorization with their figurative counterparts. Thus, *gasoline* and *beer* are both liquids, *to drink* and *to run on* are both modes of *consumption*, while *job* and *jail* can each denote a *confining, oppressive situation*. The taxonomic approach actually dates all the way back to Aristotle, but is exemplified by AI models such as those of Wilks (1978), Way (1991) and Fass (1991) and by the psychological theories of Glucksberg (1998). In approaches based on analogy, this shared representation comes in the form of isomorphic aspects of the source and target that are built from identical semantic primitives. Winston (1980), for example, as well as Carbonell (1981), Gentner (1983), Gentner *et al.* (1989), Holyoak & Thagard (1989) and Veale & Keane (1997) all propose AI models that identify a common core of structure shared by the *source* (what Davidson calls the surface meaning) and the *target* (a meaning structure of practical AI merit whose existence Davidson denies on philosophical grounds). Other symbolic approaches employ a liberal mix of the taxonomic and the analogical, from that of Hobbs (1981) to Martin (1990) to Veale (2015). Statistical approaches such as that of Mason (2004) and Shutova (2010a,b), which aim to find metaphors or to find paraphrases for metaphors that use more conventional (if not always literal) terms, avoid the thorny issue of deep meaning by remaining resolutely at the surface and by *not* building an explicit model of meaning. Nonetheless, such approaches tacitly assume the intended meaning *must* share a strong overlap with the surface meaning of its paraphrase. In the following sections we aim to craft an explicit representation for this shared meaning.

3 Persons of Interest: Good, Bad and Ugly

The *Aristotle* system of Veale & Hao (2007) is an online generator of metaphors that allows its users to select a target idea and a property to emphasize. In response, it suggests a range of source ideas that exemplify this property (such as *wolf* for *predatory*, *tiger* for *fierce*, *peacock* for *proud*, and so on), and allows the user to explore the further ramifications of any given choice (e.g. to say that one is *as predatory as a wolf* may suggest that one is *sly* and *lonely* too). *Aristotle* can also go to the Web in real time to tackle the problem of the “unknown tenor,” which is to say whenever it lacks knowledge of a target. Thus, given the metaphor *Donald Trump is a thug*, it uses its stereotypical knowledge of *thugs* to hypothesize that *Donald Trump* may be *burly*, *rude*, *brutal*, *gruff* or *sadistic*, and dispatches Web queries to quantify the validity of each hypothesis.

The current work represents a complementary approach to the problem of metaphor generation. Rather than create knowledge representations of generic animals, people or things (such as *wolves*, *thugs* and *insults*) we set out to build a knowledge base of famous proper-named individuals such as *Donald Trump*, *Tony Stark* and *Steve Jobs*. These stereotypes can then be used to construct nuanced similes, metaphors and blends, provided that we imbue our knowledge base (aka the *NOC*) with a rich enough set of *Lego*-like blocks to build its creations. Thus, as we’ll see, a system can use the *NOC* to generate XYZ metaphors such as “*Bruce Wayne is the Donald Trump of Gotham City*.”

What makes a person worthy of representation in this knowledge-base, and which aspects of each person should we actually represent? In his 1986 book *The Frenzy of Renown: Fame & its History*, Leo Braudy suggests that fame emerges from “the interplay between the common and the unique in human nature.” So the qualities that make a person worthy of figurative comparison are exactly those that seem to exist in a concentrated and quite exemplary form in that one person, yet which are so common in our experience as to be predicated of many. Our ambitions for this representation thus go much farther than that of the *Aristotle* system: we set out to represent not just the simple qualities of an iconic person (those that can be expressed as adjectives) but their many affordances as complex, fully-realized entities situated in the world, such as gender (*male* or *female*), locale (e.g. *New York*,

Gotham City), style of dress, spouses and known enemies, apt vehicles, apt weapons, domains (e.g., *science*, *politics*), taxonomic categories (e.g. *Politician*), fictive status (*real* or *fictional*), genres and creators (if fictive), typical activities (e.g., *running political campaigns*, *building casinos*), political leanings (*left*, *right* or *moderate*) and affiliations (e.g. *Eliot Ness* belongs to *The Untouchables*, *Tony Stark* to *The Avengers*, *Darth Vader* to the *Dark Side*, *Dick Cheney* to the *Republican Party*, etc.). In addition, we divide the adjectival qualities of each person in the *NOC* into their positive and negative “talking points,” the former being qualities with a positive lexical affect, such as *resolute*, *wealthy*, *agile* or *media-savvy*, the latter having an obvious negative affect, such as *evil*, *tight-fisted*, *ingratiating* or *devious*. We strive to provide at least three positive and three negative qualities for each of the 800 persons of interest in the *NOC*, so that any metaphors or blends built from these stereotypes will be more than one-note clichés, and allow for novelty, nuance and emergent inferences.

We use a conventional AI *frame* format for this knowledge: a frame for each person in the *NOC*, with its various slots and fillers as outlined above, and a frame for any of those filler concepts that are themselves worthy of further elaboration. Thus, for instance, we define frames for a person’s clothing, weapon and vehicle of choice, with slots indicating how each is to be used (e.g. one stabs with a knife, one drives a Mercedes but sails on a yacht, and one wears a hat on one’s head but wears shoes on one’s feet). One’s typical activities are associated with default locations (e.g. one *shops for shoes* in a *shopping mall* but *devises evil schemes* in an *underground lair*), while one’s taxonomic categories are themselves organized into a global ontology. A frame system such as this can be organized as a set of semantic triples (much like any semantic network), but we eschew the formalism of semantic Web technologies such as RDFS and OWL for the simplicity of spreadsheets, in which rows represent frames, columns represent slots and cells hold fillers for the corresponding frame slot. We do this to ensure that the *NOC* is easily shared and modified in an open and cumulative fashion. The *NOC* can thus be downloaded as a set of spreadsheets, containing approx. 30,000 triples in all, from the site *BestOfBotWorlds.com*. In the following section we show how the knowledge in these spreadsheets can be exploited in the service of metaphor generation.

4 Action figures in their original packaging

It may help to think of the collected personalities and paraphernalia of the NOC as a large toy-box, filled with action figures and their accessories. Our goal is to create a knowledge-base that will allow our machines the freedom to play with combinatorial possibilities much as children play with their toys. Kids show a genuinely ecumenical zeal when playing with action figures, and will happily mix characters of different genres and movie franchises (and even mix figures of very different sizes) when letting their imaginations roam freely. Kids also show little respect for orthodoxy when combining the paraphernalia of one toy genre (G.I. Joe trucks say) with the figures of another (their Harry Potter and Star Wars characters, say). We aim to give our generators of metaphors and blends the same freedom to mix and match in *their* conceptual games.

We decide on which figures to put in our figurative toy box by looking to the Web for insight, specifically to identify those personalities that are frequently used as a source of figurative comparisons. As shown in Veale (2014), the Web is replete with XYZ comparisons of the form “X is the Y of Z” where Y is a proper-named individual that is used to figuratively name a whole class of people. Examples include “red meat is the Donald Trump of cancer” (in which we infer that red meat is asserted to be an *aggressive builder* of cancers) and “the potato is the Tom Hanks of the vegetable world” (which seems to assert that potatoes are as versatile, *qua* ingredient, as Hanks is as an actor). We harvest a large body of these figurative XYZs from the Web, using the telltale capitalization of the Y field to find those that pivot around people. We observe that the use of Y’s obeys a power-law distribution, so a small number of popular Y’s hog the limelight, such as *Michael Jordan* (an icon of sporting supremacy) and *Chuck Norris* (an icon of rugged simplicity) while the bulk line up to form a long tail of names with fast-descending frequency. We take these Y’s as our starting point, and round out the list by including their various spouses, enemies, creators and so on if these are also relevant.

Though one might attempt to use information extraction techniques to automatically populate the fields of the NOC for this large set of Ys, we prefer to fill these fields manually, to build a reusable resource of the highest quality for metaphor generation. It is to this generation task that we turn next.

5 Metaphors as Memes, Tropes and Tweets

The NOC is designed to support the same level of whimsy as one might find in a child’s playground, and Twitter proves itself an ideal medium for our systems’ whimsical outputs. Limiting its outputs to just 140-character each forces a system to focus on a single central conceit and a pithy framing device. In keeping with human emanations from Twitter, these automated outputs need be neither sober nor profound, and can put the NOC to rather flip uses. Consider this simple framing device, which paints a Nietzschean maxim with colours from the NOC:

If what doesn’t kill you makes you stronger, shouldn’t being overwhelmed with ruthless ambition by [#HillaryClinton](#) make you more driven?

If what doesn’t kill you makes you stronger, shouldn’t being knocked out with an Oscar statuette by [#DanielDayLewis](#) make you more talented?

Each tweet plugs fillers from a single NOC entry into the same frame. Though no comparison is explicitly made between people, Nietzsche himself is implicitly party to the wry contrast that *is* drawn, for these tweets undercut Nietzsche’s aphorism by instantiating it in a way that is both apt *and* goofy.

Many more pieces of “stock wisdom” circulate on Twitter as empty platitudes that are just as ripe for satire using the NOC. Consider these exploits:

If clothes maketh the woman, would wearing [#HillaryClinton’s](#) pant suit make you more ambitious? Or more grasping?

Nobody’s perfect! My grandpa says to never judge a shallow diarist like [#CarrieBradshaw](#) until you’ve walked a mile in her Manolo Blahniks.

These framing devices are whimsical and fun but a system fails to get out of them anything more than it must put *into* them from the NOC. These single-scope tweets fail to leverage the protean capability of metaphor to spark new meanings at the overlap of ideas plucked from two very different domains. Yet it is a simple matter to choose a framing device that forces individuals with conflicting qualities into a juxtaposition that is at once a meaningful comparison and a sharp contrast. Consider the following tweets which again use a popular frame:

I see myself as capable, but my boss says that I make even someone as incompetent as [#EdWood](#) look like [#HillaryClinton](#).

I see myself as insightful, but my friends say that I make even someone as ignorant as [#Borat](#) look like [#AdamSmith](#).

This framing device simply chooses two figures from the NOC that are linked by an antonymous quality: *capable* for *Hillary Clinton* (and thus *inept* for *Ed Wood*) and *knowledgeable* for *Adam Smith* (and so *ignorant* for *Borat*). The NOC distribution includes a mapping of talking points to their opposites, which is largely derived from the antonym relationships found in WordNet (Fellbaum, 1998). Let's consider some XYZ metaphors in tweet-form that are also generated using the contents of NOC:

What if [#TheEmpireStrikesBack](#) were real? [#HillaryClinton](#) could be its [#PrincessLeia](#): driven yet bossy, and controversial too

What if [#TheNewTestament](#) were about [#AmericanPolitics](#)? [#MonicaLewinsky](#) could be its [#Lucifer](#): seductive yet power-hungry, and ruined too.

These XYZ framing devices simply juxtapose two figures from the NOC that share a pair of talking points, preferably one positive and one negative. In the first case, figures are chosen from the realms of the real and the fictional, to allow one to be positioned as a real-world incarnation of the other. These frames share the general linguistic structure used by [@metaphorminute](#), though the metaphors here are real, as they are predicated on a sharing of iconic qualities that the system knows to be apt.

Nonetheless, while two shared qualities anchor both of the above metaphors, each tweet actually raises *three* talking points. The additional quality – *controversial* for *Hillary/Leia* and *ruined* for *Monica/Lucifer* – is plucked at random from the talking points of the second input figure (*Hillary* or *Lucifer*) with no regard at all as to its salience to the first (*Leia* or *Monica*). Each tweet thus claims that this third quality is just as salient as the two that are known to be shared, spurring readers to reconcile this extra quality with their own knowledge. In effect, the two shared qualities act as a guarantor of the third, eliciting a good faith response from the reader. So note how the XYZ generator goes confidently out on a limb in these two metaphors:

[#NikolaTesla](#) is the real world's [#BruceBanner](#): brilliant and smart, but reclusive

What if [#BackToTheFuture](#) were real? [#NikolaTesla](#) could be its [#DocEmmettBrown](#): experimental yet nutty, and tragic too

The above tweets, each revolving around the real life figure *Nikola Tesla*, show how a third speculative quality is boldly smuggled into the target domain in plain sight, via a comparison justified by just two shared qualities. Yet in each case this extra quality can be accepted as apt, perhaps even as insightful, in its new context. Tesla was notoriously reclusive, but this can make sense for Banner too, given his fear of public carnage. Tesla also had a tragic end, in large part caused by a surfeit of ideas and a deficit of business sense. Though we cannot know how Doc Brown will end his days, a similar end seems just as conceivable, so the possibility is more thought-provoking than random. If a metaphor generator uses its knowledge in such a way as to meet its readers halfway, its readers will use their imaginations to bridge the remaining gap.

Given that two figures openly share a pair of talking points, a system may introduce a third person rather than a third quality, as in the following:

If [#MonicaLewinsky](#) is [#Lucifer](#) in a stained blue dress, who in [#TheNewTestament](#) is [#HillaryClinton](#) most like?

So the connectedness of knowledge in the NOC – essentially a dense semantic-network – allows a system to mine the same seam in multiple tweets. Here it uses the *known enemies* field of the NOC to link *Lewinsky* to *Clinton* and draw the latter into the metaphor, while also accessorizing the tweet with salient material from the former's wardrobe. Just as there can be three people in a marriage, we can also put three in a metaphor, to form a blend:

[#LexLuthor](#) combines the best of [#WarrenBuffett](#) and the worst of [#LordVoldemort](#): He is rich and successful yet also evil, hateful and bald

This two-way sharing allows anyone to be disintegrated into those facets it shares with others, so as to be *re-integrated* into an illuminating whole via blending. And though metaphors are asymmetric, we cannot help but see shades of *Lex Luthor* in the market success of *Warren Buffett*, the famed sage of Omaha, making Luthor the sage of Metropolis.

6 Strange Dreams Are Made of This

Davidson’s likening of metaphors to dreams was motivated by more than a desire to show that the “true” meaning of metaphors is just as tendentious a notion as the “true” meaning of dreams. Metaphors, like dreams, slip the surly bonds of commonsense semantics to explore areas of the imagination that challenge, even defy, literal expression. In fact, metaphor is often seen as being the engine of much of what is curious and creative in dreams, and any generative system that aims to flex its imagination should first show a mastery of metaphor.

Many dreams revolve around conflict – arising from a desire for wish-fulfillment and discord resolution – and the NOC allows a system to imagine conflict between iconic figures. Freud (1933) saw dreams as the royal road to the unconscious, so these conflicts must suggest real symbolic insights. A system must thus invent both a dream metaphor *and* its meaning, that is, its Freudian interpretation:

I dreamt I was paying kickbacks to the police chief with #MayorJoeQuimby when we were cut with a Jedi lightsaber by #LukeSkywalker

I guess #LukeSkywalker and #MayorJoeQuimby represent warring parts of my personality: the honest vs. corrupt sides. #Honest=#Corrupt

As shown above, the NOC is used to invent a conflict between two iconic figures, blending a typical activity of one with an assault with an apt weapon from another. The domain-crossed actors are chosen because they exhibit conflicting qualities (e.g. *honest* vs. *corrupt*) and it is this conflict that yields the psychological motivation – and the meaning – of a dream. Here is another pair of example tweets:

I dreamt I was holding evasive press conferences with #DonaldRumsfeld when we were pierced with a sharpened chop stick by #AungSanSuuKyi

I guess #AungSanSuuKyi and #DonaldRumsfeld represent warring parts of my personality: the caring vs. uncaring sides. #Caring=#Uncaring

For comparison, let us try and build a generator of dreams that does not use the NOC but which relies instead on a more conventional font of knowledge about dreams: the so-called *dream dictionary*. The Web (as well as countless books) is a source of the standard interpretations of a many dream symbols, and we harvest 5000 or so of these with a web spider from the popular site www.dreammoods.com.

The following pair of tweets hinges on the symbol **Reunion**: the first sets the scene, the second makes wholesale reuse of the “dictionary” interpretation:

Last night I dreamt I was a soldier, attending a reunion. What does this mean? #Reunion

Well, to dream that you are attending a reunion suggests that there are feelings from the past which you need to acknowledge and recognize.

These dreams are only partially machine-generated then, combining a scene-setting tweet that is wholly invented with a response that is human-crafted. The first tweet in each pair simply takes a triple from a knowledge-base of norms – e.g. *soldiers attend reunions* – whose agent or whose object is a match for the dream symbol. In this next example, the dream symbol matches the agent of the triple:

Last night I dreamt I was a reporter, conducting inquiries. What does this mean? #Reporter

Well, to dream that you are a newspaper reporter indicates that you are making a conscious and objective observation of your life. #Reporter

These half-machine/half-human inventions do not use the NOC, and exchange celebrity casting for a psychoanalytical insight that, though perhaps just psychobabble, carries the ring of real folk wisdom. It is an interesting question, then, as to whether we humans see greater merit in these half-human creations than in fully machine-crafted NOC creations. The world has little need of an artificial generator of ersatz dreams, of course, yet if metaphor really is the dreamwork of language as Davidson suggests, a marked preference for machine “dreams” would offer a clear vote in support of a computer’s ability to craft its own metaphors *and* meanings.

7 To Tweet, Perchance to Dream

We turn to volunteers on the crowd-sourcing platform *CrowdFlower.com* to quantify the relative merit of each kind of dream metaphor and its interpretation. We paid these volunteers a small sum to provide their ratings of the *interestingness* of each dream scenario (the first tweet in a pair), the *insightfulness* of the offered interpretation (in the second tweet of a pair), and the *symbolic richness* of a dream, regardless of whether one agrees with the interpretation or not (so, both tweets together). We sampled 30 *dream+interpretation* tweet pairs

of each kind from the outputs of their respective generators, and elicited ten ratings for each, along each of the 3 dimensions listed above. Judges were filtered for scammer-like non-engagement using the usual tests, and were asked to rate each dimension on a five-point scale (via radio buttons) ranging from *Not At All* (1) to *Very Much So* (5). The average human ratings for each dimension for each kind of dream generation is presented in Table 1.

<i>Strategy</i> → ↓ <i>Dimension</i>	<i>NOC dreams,</i> <i>interpretations</i>	<i>dream dictionary</i> <i>interpretations</i>
Interestingness	3.35 ± 0.38	2.93 ± 0.42
Insightfulness	3.22 ± 0.34	3.24 ± 0.37
Symbolism	3.12 ± 0.35	3.17 ± 0.36
All dimensions	3.23 ± 0.37	3.11 ± 0.41

Table 1. Average human ratings for dreams of both types, with their interpretations, on three dimensions (and the average of all three together).

It appears human raters judge the NOC dreams as more interesting than the corresponding non-NOC scenarios, though raters also judged interpretations for the latter (which are taken wholesale from human-crafted dream dictionaries) as ever-so-slightly more insightful than the NOC-derived alternatives. These stock psychological interpretations are also judged to make deeper use of dream symbols than their NOC counterparts. Nonetheless, only the difference in *interestingness* proves to be statistically significant (using the Wilcoxon signed-rank test), and the average of all 3 dimension taken together also favors the NOC tweets over their competition.

Ultimately, this is a small-scale evaluation of what is little more than a toy-box test of the NOC as a resource for meaningful metaphor generation. Yet we anticipate that the NOC and its many affordances – which will surely grow given the open status of the resource – will lend itself to many more uses of automated metaphor generation, in systems of varying ambition and complexity, from playful bots to quirky games to serious research.

8 Conclusions

Computational metaphor generation is a task that aims to replicate a measure of human creativity with language, and as such, it belongs squarely to the field of *Computational Creativity* (or CC; see Veale, 2012). But as a discipline, CC does not distinguish itself from other branches of AI / NLP by

its use of distinctive algorithms or particular data structures. Rather, CC is an AI sub-discipline best characterized by its philosophy, especially as it pertains to automated generation. CC systems each explore their own conceptual space of generative possibilities, as defined either by explicit axioms or by the procedural semantics implicit in their code. Though a CC system may encounter many well-formed possibilities when exploring a space, it can weigh and evaluate only a small number of these possibilities. A creative system must be more than a *mere generator* of well-formed outputs; it must be a selective producer that uses knowledge – of a domain and a target audience – to select, rank and filter possibilities, to generate products that it itself appreciates as interesting, useful and novel.

The limits of a system’s knowledge determine the limits of what it can perceive and conceive, and thus the limits of what it can appreciate in its own outputs. But as yet, just as CC has no unique algorithms or data structures to call its own, it has few knowledge sources of any real scale that can be openly shared and reused. This paper has presented a new public resource, which – though still in its early stages of development – can be shared and reused by AI researchers to build systems that squeeze novelty and value from familiar ideas. We have motivated this resource via the difficult task of *metaphor generation*, showing how AI/CC/NLP systems can use this rich symbolic knowledge to ground the meaning of their own creative outputs. Computer-generated metaphors that are insightful, interesting and resonant in their use of symbols will find many applications, from entertainment to education to the pure whimsy of the following (via a Twitterbot that makes full use of the NOC):

Hulk hate puny humans who tweet about starring in reality TV shows. Hulk smash [#KimKardashian](#) and the Rolls Royce she rode in on!

We are on more serious ground when we suggest that automatic generation will allow researchers to test metaphor theories in a constructive fashion, by building robust proofs of concept, and will allow for the creation of artificial experimental stimuli. We hope that the NOC – and other resources that build upon it – will allow researchers to explore the computational possibilities of metaphor generation with the same ease, and scale of ambition, with which they now explore automated analysis. This truly is the dream work of language research.

Acknowledgements

The author is grateful to the European Commission for funding this work via the WHIM project ([What If Machine](#); grant number 611560) and the PROSECCO project (grant number 600653). The author is also grateful for the support of Science Foundation Ireland (SFI) via the ADAPT Centre for Global Intelligent Content. and for the assistance of Alessandro Valitutti with the experiments reported in section 7. The various elements of the NOC resource can be downloaded from the Web from the website *BestOfBotWorlds.com*

References

- Carbonell, J. G. (1981). Metaphor: An inescapable phenomenon in natural language comprehension. *Report 2404*. Pittsburgh, PA. Carnegie Mellon Computer Science Dept.
- Davidson, D. (1978). What Metaphors Mean. *Critical Inquiry (Special Issue on Metaphor)*, 5(5): 31–47.
- Fass, D. (1991). met*: A Method for Discriminating Metonymy and Metaphor by Computer. *Computational Linguistics* 17(1):49–90.
- Fellbaum, C. (ed.) (1998) *WordNet: An Electronic Lexical Database* (ISBN: 0-262-06197-X). Cambridge, MA: MIT Press.
- Freud, S. (1933). *New Introductory Lectures on Psychoanalysis*. London, UK: Penguin Freud Library 2.
- Gentner, D., Falkenhainer, B. and Skorstad, J. (1989). Metaphor: The Good, The Bad and the Ugly. In *Theoretical Issues in NLP*, Yorick Wilks (Ed.) Hillsdale, NJ: Lawrence Erlbaum Associates.
- Glucksberg, S. (1998). Understanding metaphors. *Current Directions in Psychological Science*, 7:39-43.
- Hobbs, J. (1981). Metaphor Interpretation As Selective Inferencing. In *Proceedings of IJCAI'81, the 7th International Joint Conference on Artificial Intelligence*, Vancouver, BC, Canada, vol. 1:85–91.
- Holyoak, K. and Thagard, P. (1989). Analogical Mapping by Constraint Satisfaction. *Cognitive Science* 13: 295–355.
- Martin, J. H. (1990). *A Computational Model of Metaphor Interpretation*. New York: Academic Press.
- Mason, Z. J. (2004). CorMet: A Computational, Corpus-Based Conventional Metaphor Extraction System, *Computational Linguistics*, 30(1):23–44.
- Shutova, E. (2010a). Models of metaphor in NLP. In *the Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics, Uppsala, Sweden, 688–697*.
- Shutova, E. (2010b). Metaphor Identification Using Verb and Noun Clustering. In *the Proc. of the 23rd COLING, the International Conference on Computational Linguistics*.
- Veale, T. and Keane, M.T. (1997). The Competence of Sub-Optimal Structure Mapping on ‘Hard’ Analogies. In *Proceedings of IJCAI'97, the 15th International Joint Conference on Artificial Intelligence*, Nagoya, Japan.
- Veale, T. and Hao, Y. (2007). Comprehending and Generating Apt Metaphors: A Web-driven, Case-based Approach to Figurative Language. In *Proc. of the 22nd AAAI conference of the Association for the Advancement of Artificial Intelligence*.
- Veale, T. (2012). *Exploding the Creativity Myth: The Computational Foundations of Linguistic Creativity*. London, UK: Bloomsbury/Continuum.
- Veale, T. (2014). The ABCs of XYZs: Creativity and Conservativity in Humorous Epithets. In: *J. Manjaly & B. Indurkha (Eds.) Cognition, Experience, and Creativity*. New Delhi: Orient Blackswan.
- Veale, T. (2015). Unnatural Selection: Seeing Human Intelligence in Artificial Creations. *Journal of General Artificial Intelligence* (special issue on Computational Creativity, Concept Invention, and General Intelligence), 6(1):5-20.
- Veale, T., Valitutti, A. and Li, G. (2015). Twitter: The Best of Bot Worlds for Automated Wit. In *Proceedings of HCII-2015, the 17th International Conference on Human-Computer Interaction (Distributed, Ambient and Pervasive Interactions)*, Los Angeles, California, 2-7.
- Way, E. C. (1991). *Knowledge Representation and Metaphor. Studies in Cognitive systems*. Kluwer Academic, Amsterdam.
- Wilks, Y. (1978). Making preferences more active. *Artificial Intelligence*, 11(13): 197–223.
- Winston, P. (1980). Learning and reasoning by analogy. *Communications of the ACM*, 23(12): 689-703.